

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Influence of partner behaviour on overspecification

Permalink

<https://escholarship.org/uc/item/4pv876q5>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Loy, Jia E.

Smith, Kenny

Publication Date

2020

Peer reviewed

Influence of partner behaviour on overspecification

Jia E. Loy (jia.loy@ed.ac.uk)

Linguistics and English Language, University of Edinburgh, United Kingdom

Kenny Smith (kenny.smith@ed.ac.uk)

Linguistics and English Language, University of Edinburgh, United Kingdom

Abstract

Speakers often overspecify by using colour adjectives redundantly in referential communication. We investigated whether this tendency to overspecify is influenced by a partner's linguistic behaviour, and whether the effect is enhanced by lexical repetition and semantic relatedness. We used a director-matcher task in which speakers interacted with either a consistently overspecific or a consistently optimal partner. Our results show that partner behaviour influences overspecification. An analysis over time indicates that speakers tended to overspecify at the outset, but reduced this behaviour over interaction with an optimal partner much more than with an overspecific partner. This may suggest that overspecification (at least with colour modifiers) is the "default" behaviour, with speakers adapting to optimality in a partner's linguistic behaviour.

Keywords: overspecification; partner alignment; referential communication; pragmatics

Introduction

When referring to an object, speakers may encode information that is superfluous and not required for a listener to uniquely identify the object being described in context (e.g. requesting the "blue sock" in a context where only one sock is present). This behaviour is known as overspecification and is a recurring theme in research on the pragmatic processes underlying reference production and comprehension. Studies on referential communication often note that overspecification is unexpected given the Gricean Maxim of Quantity, which states that speakers should provide as much information as necessary but no more (Grice, 1975). This raises the question as to what causes speakers to produce overspecific object references. Here, we examine the role of an interlocutor's linguistic behaviour on a speaker's tendency to overspecify.

Previous research attributes referential overspecification to a variety of factors. These may be speaker-centric, such as an attempt to facilitate production before visual processing is complete (Pechmann, 1989), or a heuristic tendency to encode attributes that are perceptually salient (Koolen, Gatt, van Gompel, Krahmer, & Van Deemter, 2016). Speakers tend to overspecify colour more than, for example, size (Koolen, Gatt, et al., 2016), a finding which has been attributed to the relative salience and absoluteness of an object's colour (Tarenskeen, Broersma, & Geurts, 2015). The rate of overspecification is also sensitive to visual context: speakers overspecify more with more complex scenes, e.g. when distractors vary in number or colour (Koolen, Goudbeek, & Krahmer, 2013; Koolen, Krahmer, & Swerts, 2016). These find-

ings imply a general link between cognitive demand on a speaker and their tendency to overspecify.

Other researchers argue for overspecification as a listener-oriented process. Experimental evidence suggests that speakers may overspecify in order to facilitate their interlocutor's search for an intended referent (Arts, Maes, Noordman, & Jansen, 2011; Rubio-Fernández, 2016). Speakers in Rubio-Fernández (2016) overspecified only when target objects were atypically coloured (e.g. a pink banana), and were more likely to do so when instructions emphasised the possibility of communication breakdown with the listener. These findings are consistent with evidence from comprehension to suggest that listeners expect speakers to use prenominal adjectives rationally: on hearing a scalar adjective such as "tall ...", for instance, listeners begin to narrow their visual search down to a set of contrastive objects (e.g. a tall and a short glass; Sedivy, Tanenhaus, Chambers, & Carlson, 1999). Listener-oriented overspecification has also been observed in speakers' choice of lexical names: directors who had entrained with a matcher on an overspecific subordinate name (e.g. "pennyloafer" for a shoe) stopped overspecifying and reverted to simpler, basic-level terms such as "shoe" with a new matcher (Brennan & Clark, 1996). Together, these findings suggest that overspecification is interlocutor-dependent, and support a growing literature to show that speakers take into account their partner's perspective in various aspects of reference production (e.g. Hanna & Tanenhaus, 2004; Chambers, Tanenhaus, & Magnuson, 2004; Brennan & Hanna, 2009).

Less is known, however, about the role of a partner's linguistic behaviour on a speaker's production of overspecific expressions. The majority of studies on referential overspecification have focussed on one-sided communication, either with hypothetical listeners or fixed speaker/addressee roles. However, there is reason to believe that a speaker's tendency to overspecify may be influenced by their interlocutor's referential behaviour in bi-directional interaction.

Referential communication studies involving dialogue show that interlocutors often adopt their partner's forms of expression over time, leading to coordinated behaviour at many levels of linguistic structure (e.g. Branigan, Pickering, & Cleland, 2000; Cleland & Pickering, 2003; Branigan, Pickering, Pearson, McLean, & Brown, 2011). Pickering and Garrod (2004) explain this via an Interactive Alignment Model, which proposes that conversation partners come to align on

their linguistic representations through dialogue, leading to a tendency to repeat each other's choices at various levels, from low-level phonetic realisations to high-level situational models. Branigan et al. (2000) provide a demonstration of such alignment at the syntactic level. They showed that prepositional object (PO) prime descriptions (e.g. "the pirate giving the banana to the doctor") provided by a confederate were more likely to elicit PO target descriptions from the participant, while double object (DO) prime descriptions (e.g. "the pirate giving the doctor the banana") were more likely to elicit DO target descriptions. This alignment process is assumed to be largely automatic and subconscious (Pickering & Branigan, 1998), with percolation of alignment between levels enhancing the effect (Garrod & Pickering, 2004). For instance, syntactic alignment is enhanced by lexical repetition between prime and target (e.g. *give-give* vs. *give-show*; Branigan et al., 2000), as well as by semantic relatedness between prime and target (e.g. *sheep-goat* vs. *sheep-knife*; Cleland & Pickering, 2003). Such alignment has been observed even when a partner's choice of expression may be more effortful for a speaker to produce, for instance with less frequent lexical names such as "coach" for a bus (Branigan et al., 2011), or dispreferred modifier attributes such as orientation as opposed to colour of an object (Goudbeek & Krahmer, 2012, experiment1).

Goudbeek and Krahmer (2012) also addressed the question of alignment in the production of overspecific referring expressions. In their experiment 3, they showed that participants who heard overspecific prime descriptions containing a preferred and a dispreferred attribute (e.g. "the red chair seen from the front") when either attribute would suffice were more likely to produce similarly overspecific target descriptions compared to participants who heard primes with only one attribute. However their study only examined alignment occurring within semantic categories (e.g. furniture objects), and relied on interaction with a hypothetical interlocutor.

In the current experiment, we attempt to more closely approximate authentic interaction by presenting the experiment as a communicative game with a remote partner in real-time. We test whether speakers align with a partner's linguistic behaviour to produce overspecific object descriptions using prenominal colour adjectives, comparing interaction with a partner who consistently overspecifies with one who is consistently optimal. In addition, we test whether this alignment is enhanced by lexical repetition between prime and target modifiers (e.g. "red sock" to "red glove") as well as semantic overlap between prime and target categories (e.g. sock to glove vs. sock to cup). Previous work examining lexical and semantic boost effects in alignment have mainly concentrated on speakers' choice of syntactic structure (e.g. Cleland & Pickering, 2003; Hartsuiker, Bernolet, Schoonbaert, Speybroeck, & Vanderelst, 2008). These studies suggest that syntactic encoding is affected by repetition of lexical and conceptual information. Referential production experiments have also shown that speakers' choice of modifier (size vs. colour)

is influenced by repetition of semantic content in the form of the noun being described (Heller & Chambers, 2004). Less is known, however, about whether the pragmatic phenomenon of overspecification, i.e. whether or not a redundant modifier is used, is similarly sensitive to lexical and semantic effects.

Method

The experiment was a director-matcher task in which participants alternated between describing and matching pictures. On critical trials, participants heard a partner-produced prime description which either contained an overspecific colour modifier (overspecific partner) or featured only the bare noun (optimal partner). On the following target trial, participants produced an object description for their partner, where colour was never necessary to distinguish between objects in the display. We manipulated whether the prime and target objects have the same or different colours, and whether they were drawn from the same or different semantic categories.

Participants

Sixty-eight self-reported native speakers of English between the ages of 18 and 35 were each paid £6 to take part.

Materials

Critical stimuli We created four different coloured image variants (red, blue, green, yellow) of each of eight common nouns from two semantic categories (clothing: cap, glove, scarf, sock; kitchenware: bowl, fork, mug, pot). Each image was associated with a voice recording specifying the noun item and colour modifier, e.g. "*the red cap*". The recordings were produced by a male speaker of British English in a single recording session. For each object, the same token of each colour modifier was cross-spliced onto the bare noun to create the colour-modified version.

Filler stimuli The experiment also included two types of filler stimuli: images of easily recognisable natural objects, and photographs of human facial expressions. For natural object fillers, we selected 32 images (16 animals; 16 fruits and vegetables) with high nameability (name agreement $\geq 90\%$ and H value ≤ 0.5) from the Bank of Standardised stimuli (BOSS)¹. Each image was associated with a recording of its modal name as determined through the study's norming procedure, e.g. "*the giraffe*". For facial expression fillers, eight photographs (four men; four women) for each of six emotions (angry, disgusted, frightened, happy, sad, surprised) were taken from the Karolinska Directed Emotional Faces database (KDEF)². These were converted to grayscale to discourage participants from relying on colour-modified facial features such as hair and eye colour in their descriptions. Each photo

¹BOSS is a set of normative colour images created for cognitive and psycholinguistic research (Brodeur, Dionne-Dostie, Montreuil, & Lepage, 2010)

²KDEF is a collection of purposely-designed photos of expressions developed for psychological research (Lundqvist, Flykt, & Ohman, 1998)

was associated with a recording describing the subject’s facial expression. A variety of descriptions were used, e.g. “*the woman who looks angry*”, “*the angry-looking woman*”, “*the angry one*”. All filler recordings were produced by the same speaker who recorded the critical stimuli.

Design

The experiment utilised a 2 x 2 within-subjects design manipulating colour (same vs. different for prime and target) and semantic category (within vs. across category between prime and target). Additionally, partner behaviour (overspecific vs. optimal) was manipulated between-subjects.

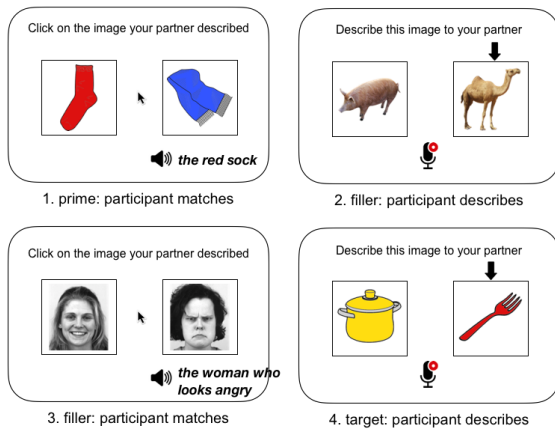


Figure 1: Example of a four-trial sequence in a critical item from the same colour–across category condition.

The experiment included 32 critical items—eight per within-subjects condition. Each critical item consisted of a four-trial sequence: a prime (participant matching), a filler (participant describing), a filler (participant matching), and a target (participant describing) trial. Within each condition, each of the eight noun items (four per category) appeared once as a prime and once as a target referent, with the four objects from each category appearing once in each colour as prime and once as target.

On critical prime and target trials, the display featured a target image alongside a distractor, chosen randomly from the full set of objects with the constraint that it differed from the target in both noun and colour (i.e. colour was never necessary to distinguish between the two objects). Filler trials similarly presented two images, both either natural objects or facial expressions. These were included to reduce the connection between the prime and target trials (cf. Goudbeek & Kraemer, 2012). Fig. 1 shows an example of four trials that constitute a critical item. The relative positions of target and distractor images were randomised, with the target appearing equally on each side in each condition.

The experiment included an additional 128 filler trials: 48 natural object trials, 48 facial expression trials, and 32 trials featuring the same images used in critical trials. For natural object fillers, the target and distractor were always from the

same category (i.e. animals or fruits and vegetables) to discourage participants from relying on superordinate category-based naming strategies. These fillers were designed to be easily nameable in order to elicit unmodified referring expressions. For facial expression fillers, the target and distractor always depicted subjects of the same gender to encourage participants to focus on facial expression in their descriptions. These fillers were included primarily as a more open-ended description task, to distract participants from the actual focus of the experiment. For the 32 filler trials featuring the same images as critical trials, half were displays which required colour for disambiguation (e.g. a red sock and a blue sock), while the other half were displays which required the object name for disambiguation (e.g. a red sock and a red ball). These were included to prevent participants from learning to rely on specifying only colour or only noun over the course of the experiment.³ The overspecific partner always specified colour on these fillers regardless of whether it was necessary, while the optimal partner only specified colour when it was necessary (i.e. the same item in different colours). For all three filler types, half of the trials were participant-describing trials and the other half participant-matching trials.

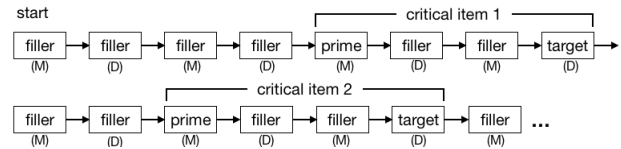


Figure 2: Example timeline of trial order. M and D denote whether the participant is matching (M) or describing (D).

The order of presentation of trials was randomised on each run with the constraints that at least four filler trials preceded the first critical item, and at least two filler trials occurred between critical items (see fig. 2).

Procedure

Participants were tested individually in sound-attenuated booths. Prior to the experiment, participants were told they would be playing a picture description and matching game over a networked connection with a participant located at a partner university.

After clicking begin, a “connecting” screen was displayed for 15 s to simulate a wait for their partner to connect to the network. The experiment always began with a participant-matching trial, and alternated between description and match trials. On description trials, the target and distractor images appeared side-by-side with an arrow pointing to the target. After 500 ms, a microphone symbol below turned red to signify that participants were being recorded. Participants were instructed to click on the microphone when they had finished speaking to send their description to their partner. A “wait” message appeared for a set delay to simulate their partner

³This type of filler never appeared within critical items to avoid directly influencing speakers before a target trial.

selecting an object. This was fixed at 2,000 ms for critical trials, and randomly variable between 1,800 and 3,000 ms on filler trials, with the delay decreasing as the experiment progressed. On match trials, the target and distractor images appeared side-by-side. Playback of the audio recording associated with the target began after a delay, fixed at 2,000 ms on critical trials, and randomly variable between 1,800 and 3,000 ms on filler trials, with the delay decreasing as the experiment progressed. After this, the mouse pointer appeared at the centre, along with the instruction text “Click on the picture your partner described” below the images. No feedback was provided after the participant clicked on an image.

To simulate more naturalistic interaction, the partner audio included a variety of disfluencies (filled pauses, repetitions, prolongations). The first three partner descriptions were disfluent, after which approximately 10% of fillers were disfluent, with the probability of a disfluency occurring being higher on facial expression fillers, and decreasing as the experiment progressed.

At the end of the experiment participants were debriefed, during which the partner manipulation was revealed. Participants were then verbally asked whether they had suspected that they were not interacting with a partner in real-time; only participants who explicitly confirmed that they believed the interaction was real were included in our analyses.

Results

Data and analysis

We excluded data from 20 participants who indicated during debrief suspicion about the nature of the interaction. Thus, the final dataset we analysed consisted of 48 participants (24 per partner condition).⁴

We coded participants’ descriptions for overspecification. This was defined as including both the colour modifier and noun item in the description. Across 1,536 critical trials, 977 overspecific descriptions were recorded: 644 with the overspecific partner and 333 with the optimal partner.

We analysed the outcome variable of whether or not participants produced an overspecific description using logistic mixed effects regression. To analyse the overall influence of their partner’s behaviour, we modelled the outcome variable by the fixed effects of target colour (same/different), semantic category (within/across), and partner (optimal/overspecific), with all predictors sum-coded. Participant and target object intercepts, and by-participant slopes for target colour and semantic category were added as random effects. To test whether participants were influenced by their partner’s behaviour over time, we conducted a second analysis including critical trial progression as a predictor. The model structure was the same as before, with the addition of the fixed effect of trial (coded such that the intercept was centred on the first critical trial).

⁴We note that analysis on the full dataset including all participants shows approximately the same pattern of results. We report these results in footnotes 6 and 7.

Effect of colour, category and partner on overspecification behaviour

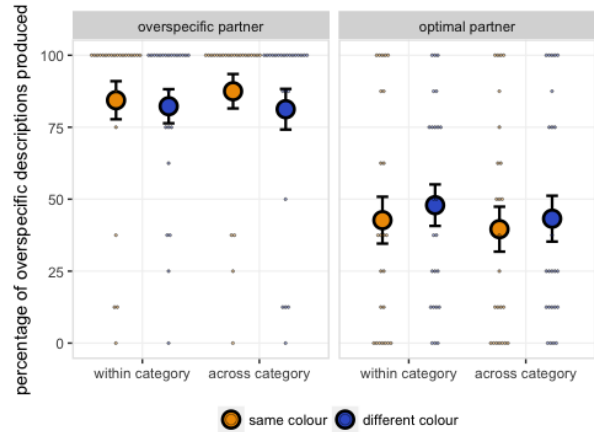


Figure 3: Mean percentages of overspecific descriptions produced by participants in each condition. Error bars represent ± 1 standard error of by-participant means. Dots represent individual participant means.

Fig. 3 shows the overall mean percentage of overspecific descriptions produced by participants in each condition. The analysis showed a main effect of partner: Participants were more likely to overspecify with an overspecific partner than an optimal partner, $\beta = 4.94$, $SE = 1.15$, $p < .001$; as well as a marginal main effect of colour: Participants were marginally less likely to overspecify when the prime and target objects differed in colour $\beta = -0.73$, $SE = 0.39$, $p = .06$. There was also a partner by colour interaction: Participants interacting with an overspecific partner were more likely to overspecify when prime and target objects shared the same colour compared to participants interacting with an optimal partner, $\beta = -2.21$, $SE = 0.77$, $p = .004$. This likely reflects a lexical boost effect of priming that occurred with an overspecific partner.⁵ There was no effect of semantic category nor its interaction with any of the other predictors.⁶

Influence of partner’s behaviour over time

Because we were interested in whether participants’ rate of overspecification would change over the course of interacting with their partner, we conducted a second analysis that looks at overspecification across trials. Fig. 4 shows the mean percentage of overspecific descriptions produced by participants in each condition over the course of the experiment’s

⁵Although Fig. 3 appears to indicate that participants interacting with an optimal partner were more likely to overspecify when prime and target objects differed in colour, separate analyses by partner condition showed an effect of colour only in the overspecific partner condition and no difference in the optimal partner condition. This suggests that the interaction was driven mainly by a difference by colour with the overspecific partner.

⁶Analysis including all participants: Partner effect, $\beta = 4.38$, $SE = 0.86$, $p < .001$; partner:colour interaction, $\beta = -1.23$, $SE = 0.47$, $p = .008$; colour effect *n.s.*

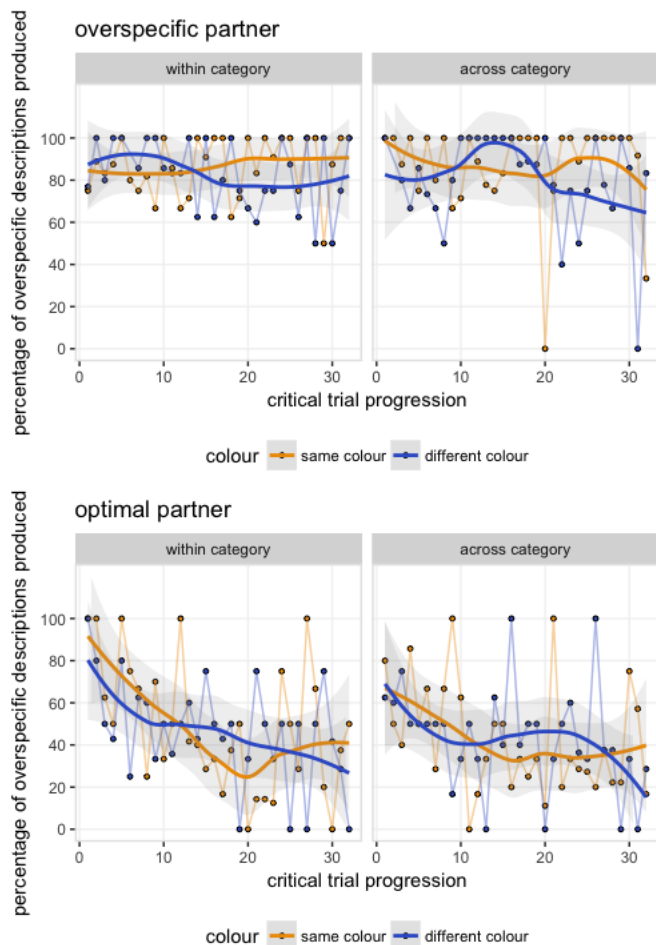


Figure 4: Mean percentages of overspecific descriptions produced by participants over critical trial progression. The curved lines represent the data with a loess smoothing curve fitted. The grey ribbons represent 95% confidence intervals around the smooth.

progress. There were effects of partner, with more overspecification occurring with an overspecific partner, $\beta = 4.47$, $SE = 1.36$, $p = .001$; of colour, with less overspecification when prime and target objects differed in colour, $\beta = -2.29$, $SE = 0.71$, $p = .001$; and a partner by colour interaction, $\beta = -3.25$, $SE = 1.21$, $p = .007$, reflecting a lexical boost effect of priming with an overspecific partner. Turning to the rate of overspecification over time, there was an effect of trial progression, $\beta = -0.73$, $SE = 0.13$, $p < .001$, indicating that, overall, participants were less likely to overspecify as the experiment progressed. Trial progression interacted with colour, $\beta = 0.77$, $SE = 0.27$, $p = .004$, reflecting a greater decrease in overspecification when prime and target differed in colour. Importantly, trial progression interacted with partner, $\beta = 0.63$, $SE = 0.25$, $p = .01$, reflecting a greater decrease in overspecification when interacting with an optimal partner. This difference is evident in the rates of overspecification over trials in Figure 4, and was confirmed by separate anal-

yses on the data for each partner condition, which showed a larger effect of trial progression in the optimal partner condition, $\beta = -1.00$, $SE = 0.14$, $p < .001$, than in the overspecific partner condition, $\beta = -0.60$, $SE = 0.28$, $p = .03$.⁷

Discussion

We investigated whether a speaker's tendency to overspecify is influenced by their partner's linguistic behaviour, and whether this effect is enhanced by lexical repetition and semantic relatedness. We found an overall influence of partner behaviour, such that speakers were more likely to produce overspecific colour modifiers with a partner who consistently overspecified compared to one who was consistently optimal. These results are consistent with Goudbeek and Krahmer (2012), who found that speakers were influenced by their partner to overspecify in their use of attributes to describe objects. More generally, they demonstrate that, as with many other aspects of linguistic behaviour, overspecification is susceptible to the tendency to align with a partner, strengthening the evidence for interactive alignment in dialogue.

As expected, the alignment we observed was enhanced by lexical overlap between prime and target descriptions: speakers interacting with the overspecific partner were more likely to overspecify when prime and target objects shared the same colour. Previous work on lexical boost effects have mainly focussed on syntactic alignment between speakers; however, our results show that lexical enhancement extends to alignment on a pragmatic aspect of production—overspecification. The fact that overspecification with the optimal partner was not similarly affected by colour repetition is unsurprising, since the optimal partner's prime descriptions never included a colour modifier.

Our results indicated no effect of category on speakers' likelihood of overspecifying with either partner. This suggests that alignment in overspecification is not affected by semantic relatedness between prime and target. This appears at odds with Cleland and Pickering (2003), who found that speakers were more likely to repeat a partner's syntactic structure when the prime and target contained nouns from the same semantic category. Cleland and Pickering attribute the semantic boost to increased activation of related concepts by the prime noun (e.g. *sheep* activates *goat* more than it does *knife*). It is possible, however, that the process of syntactic planning activates semantically related concepts in a way that pragmatic encoding does not. Another simpler explanation is that the semantic categories we used (clothes, kitchenware) were not distinct enough for a within-category advantage to be observed. This seems possible particularly when considering that the category distinctions in Cleland and Pickering (2003) were frequently animate/inanimate objects. Pre-

⁷ Analysis including all participants: Partner effect, $\beta = 3.98$, $SE = 0.97$, $p < .001$; colour effect, $\beta = -1.02$, $SE = 0.42$, $p = 0.02$; partner:colour interaction, $\beta = -1.97$, $SE = 0.74$, $p < .01$; trial effect, $\beta = -0.67$, $SE = 0.09$, $p < .001$; trial:colour interaction, $\beta = 0.43$, $SE = 0.19$, $p = .02$; trial:partner interaction, $\beta = 0.44$, $SE = 0.19$, $p = .02$

vious research also highlights the role of colour pertinence on colour overspecification, with objects for which colour is a central property (e.g. clothes, shoes) more susceptible to overspecification (Rubio-Fernández, 2016). A category effect on speakers' tendency to align may thus be observed with more distinct categories such as items of clothing and geometric shapes (cf. Rubio-Fernández, 2019). Further work examining semantic enhancement using a range of category distinctions would provide a more concrete picture of the role of semantic relatedness in alignment.

Turning to the influence of partner behaviour over time, our results indicate that speakers adapted differently to an overspecific and an optimal partner. Unexpectedly, rather than increase their rate of overspecification with an overspecific partner, speakers in both partner conditions tended to overspecify at the outset and decrease this behaviour over time. Notably, those in the optimal partner condition decreased their rate of overspecification much more dramatically, likely in response to their partner's linguistic behaviour. In other words, speakers appeared to be more influenced by the optimal rather than the overspecific partner's behaviour. Anecdotally, we note that this is consistent with reports from several participants during debrief that they stopped producing redundant colour adjectives when they noticed their partner did not produce them. The fact that overspecification seemed to be the 'default' manner of expression initially for speakers in both conditions suggests that this behaviour might be due to speaker-internal processes. This supports the view held by some psycholinguists that overspecification, at least with colour mention in reference production, may be the result of an intrinsic tendency to encode information that is visually salient or easily cognisable (e.g. Pechmann, 1989). This could explain why speakers in our experiment started out overspecifying colour, but were more likely to cease doing so over the course of interacting with an optimal partner.

It should be noted, however, that our results do not rule out the possibility that speakers would also align with an overspecific partner over time. The strong tendency in speakers to overspecify from the outset limited our opportunity to observe any increase as the experiment progressed. This tendency was likely a result of the fact that we targetted colour overspecification (as opposed to say size or pattern; Tarensken et al., 2015; Belke & Meyer, 2002), and additionally used prototypical colours that were visually distinct (cf. Viethen, Goudbeek, & Krahmer, 2016) and objects which tend to elicit colour-modified expressions (Rubio-Fernández, 2016). It is plausible that speakers may be similarly influenced by an overspecific partner in a context where they are less likely to overspecify by default, for instance with objects that differ on a scalar dimension such as size or number. Another possibility would be to explore the effect of interacting with a partner who modifies their linguistic behaviour mid-interaction—we might see alignment with overspecification emerge as a partner shifts from being consistently optimal to being overspecific.

Finally, we note that the paradigm we employed still somewhat departs from authentic dialogue. Although we designed our recordings to simulate realistic partner behaviour, and only analysed participants who reported having believed the partner manipulation, such methods may not be fully comparable with actual interaction between interlocutors. Notably, however, a recent study by Out, Goudbeek, and Krahmer (2020) replicated Goudbeek and Krahmer's (2012) finding of alignment with a partner's choice of modifier (colour vs. orientation) in a more naturalistic dialogue setting. This highlights the robustness of alignment behaviour, at least with modifier encoding, in authentic interactive contexts. The alignment of overspecific modifier use in authentic interaction would be a useful avenue for future research to pursue.

Conclusion

We investigated the influence of a partner's linguistic behaviour on speakers' tendency to produce overspecific colour modifiers, and whether this effect would be enhanced by lexical repetition and semantic relatedness. Speakers were influenced by their partner's behaviour to produce more overspecific referential descriptions with an overspecific partner compared to an optimal partner. This effect was magnified by lexical repetition when the prime and target shared the same colour. The behaviour of speakers over time suggests that they tended to overspecify at the outset, but were influenced by an optimal partner more than an overspecific partner to reduce this behaviour over interaction. This suggests that speakers in our experiment were adapting to optimality rather than overspecification in their partner's linguistic behaviour.

Acknowledgments

This research received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 681942), held by KS. The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

- Arts, A., Maes, A., Noordman, L., & Jansen, C. (2011). Overspecification facilitates object identification. *Journal of Pragmatics*, 43(1), 361–374.
- Belke, E., & Meyer, A. S. (2002). Tracking the time course of multidimensional stimulus discrimination: Analyses of viewing patterns and processing times during “same”-“different” decisions. *European Journal of Cognitive Psychology*, 14(2), 237–266.
- Branigan, H. P., Pickering, M. J., & Cleland, A. A. (2000). Syntactic co-ordination in dialogue. *Cognition*, 75(2), B13–B25.
- Branigan, H. P., Pickering, M. J., Pearson, J., McLean, J. F., & Brown, A. (2011). The role of beliefs in lexical alignment: Evidence from dialogs with humans and computers. *Cognition*, 121(1), 41–57.

- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1482–1493.
- Brennan, S. E., & Hanna, J. E. (2009). Partner-specific adaptation in dialog. *Topics in Cognitive Science*, 1(2), 274–291.
- Brodeur, M. B., Dionne-Dostie, E., Montreuil, T., & Lepage, M. (2010). The Bank of Standardized Stimuli (BOSS), a new set of 480 normative photos of objects to be used as visual stimuli in cognitive research. *PloS one*, 5(5), e10773.
- Chambers, C. G., Tanenhaus, M. K., & Magnuson, J. S. (2004). Actions and affordances in syntactic ambiguity resolution. *Journal of experimental psychology: Learning, memory, and cognition*, 30(3), 687–696.
- Cleland, A. A., & Pickering, M. J. (2003). The use of lexical and syntactic information in language production: Evidence from the priming of noun-phrase structure. *Journal of Memory and Language*, 49(2), 214–230.
- Garrod, S., & Pickering, M. J. (2004). Why is conversation so easy? *Trends in Cognitive Sciences*, 8, 8–11.
- Goudbeek, M., & Krahmer, E. (2012). Alignment in interactive reference production: Content planning, modifier ordering, and referential overspecification. *Topics in Cognitive Science*, 4(2), 269–289.
- Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41–58). Brill.
- Hanna, J. E., & Tanenhaus, M. K. (2004). Pragmatic effects on reference resolution in a collaborative task: Evidence from eye movements. *Cognitive Science*, 28(1), 105–115.
- Hartsuiker, R. J., Bernolet, S., Schoonbaert, S., Speybroeck, S., & Vanderelst, D. (2008). Syntactic priming persists while the lexical boost decays: Evidence from written and spoken dialogue. *Journal of Memory and Language*, 58(2), 214–238.
- Heller, D., & Chambers, C. G. (2004). Would a blue kite by any other name be just as blue? effects of descriptive choices on subsequent referential behavior. *Journal of Memory and Language*, 70, 53–67.
- Koolen, R., Gatt, A., van Gompel, R. P., Krahmer, E., & Van Deemter, K. (2016). Viewing time affects overspecification: Evidence for two strategies of attribute selection during reference production. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 159–164).
- Koolen, R., Goudbeek, M., & Krahmer, E. (2013). The effect of scene variation on the redundant use of color in definite reference. *Cognitive Science*, 37(2), 395–411.
- Koolen, R., Krahmer, E., & Swerts, M. (2016). How distractor objects trigger referential overspecification: testing the effects of visual clutter and distractor distance. *Cognitive science*, 40(7), 1617–1647.
- Lundqvist, E., Flykt, A., & Ohman, A. (1998). *The Karolinska Directed Emotional Faces - KDEF*. CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet, ISBN 91-630-7164-9.
- Out, C., Goudbeek, M., & Krahmer, E. (2020). Do speaker's emotions influence their language production? studying the influence of disgust and amusement on alignment in interactive reference. *Language Sciences*, 78, 101255.
- Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics*, 27(1), 89–110.
- Pickering, M., & Branigan, H. P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language*, 39(4), 633–651.
- Pickering, M., & Garrod, S. (2004). Towards a mechanistic psychology of dialogue. *Behavioural and Brain Sciences*, 27, 169–226.
- Rubio-Fernández, P. (2016). How redundant are redundant color adjectives? an efficiency-based analysis of color overspecification. *Frontiers in Psychology*, 7, 153.
- Rubio-Fernández, P. (2019). Overinformative speakers are cooperative: Revisiting the Gricean Maxim of Quantity. *Cognitive Science*, 43(11), e12797.
- Sedivy, J. C., Tanenhaus, M. K., Chambers, C. G., & Carlson, G. N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, 71(2), 109–147.
- Tarenskeen, S., Broersma, M., & Geurts, B. (2015). Overspecification of color, pattern, and size: salience, absoluteness, and consistency. *Frontiers in Psychology*, 6, 1703.
- Viethen, J., Goudbeek, M., & Krahmer, E. (2016). The impact of colour difference and colour codability on reference production. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 1084–1089).